arXiv:1602.01792v3 [cs.IR] 16 Jun 2016

Random Forest DBSCAN for USPTO Inventor Name Disambiguation*

Kunho Kim[‡], Madian Khabsa[§], C. Lee Giles[†‡]

[‡]Computer Science and Engineering
[†]Information Sciences and Technology
The Pennsylvania State University
University Park, PA 16802, USA

[§]Microsoft Research
One Microsoft Way
Redmond, WA 98005, USA

kunho@cse.psu.edu, madian.khabsa@microsoft.com, giles@ist.psu.edu

Abstract

Name disambiguation and the subsequent name conflation are essential for the correct processing of person name queries in a digital library or other database. It distinguishes each unique person from all other records in the database. We study inventor name disambiguation for a patent database using methods and features from earlier work on author name disambiguation and propose a feature set appropriate for a patent database. A random forest was selected for the pairwise linking classifier since they outperform Naive Bayes, Logistic Regression, Support Vector Machines (SVM), Conditional Inference Tree, and Decision Trees. Blocking size, very important for scaling, was selected based on experiments that determined feature importance and accuracy. The DBSCAN algorithm is used for clustering records, using a distance function derived from random forest classifier. For additional scalability clustering was parallelized. Tests on the USPTO patent database show that our method successfully disambiguated 12 million inventor mentions within 6.5 hours. Evaluation on datasets from USPTO PatentsView inventor name disambiguation competition shows our algorithm outperforms all algorithms in the competition.

1 Introduction

One of the most frequent queries for digital library search system is a person name. An example is to find all relevant records of a particular person. For a patent database, users may want to find the list of patents of a certain inventor. This query can be problematic if there is no unique identifier for each person. In that case, a method must be used to distinguish between person records in

*A shorter version of this paper was published in JCDL [14]

the database. This is often referred to as the personal name disambiguation problem.

There are several factors that make this problem hard. Firstly, there can be several different formats for displaying one person's name. For example, one record has the full name "John Doe", while another contains only the iniital of first name, "J. Doe". More importantly, there are some common names that many people share. We can see this problem often with Asian names. Statistics from Wikipedia[1] show that 84.8% of the population have one of the top 100 popular surnames in China, while only 16.4% of common names in United States. For some records the first and last name is reversed, especially for certain groups that put the last name first for their full name. Lastly, typographical errors and foreign characters can also challenge disambiguation. In addition, because of the large number of records, usually millions, manual disambiguation of all records is not feasible (which is even then not perfect) and automated methods have to be used. An automatic name disambiguation algorithm typically consists of two parts. The first is a pairwise linkage classifier that determines whether each pair of records are from the same person or not [25]. The second is a clustering algorithm, grouping records for each unique person using the classifier.

Here, we propose to use an author name disambiguation algorithm for the patent inventor database. Our algorithm follows the typical steps of author name disambiguation, but with a newly proposed set of features from patent metadata. Having experimented with several different classifers, we use a random forest classifier to train for pairwise linkage classification and use DBSCAN for clustering for disambiguation. We use the publicly available USPTO database for testing. Recently there was an inventor name disambiguation competition for this database. Raw data is publicly available via the competition's web page[2]. This raw data contains all published US patent grants from 1976 to 2014. Although we didn't participate in the competition, we used the same training and test datasets used in the competition for evaluation. The competition's evaluation results show our algorithm to be superior to other suggested algorithms in the competition. A detailed explanation of dataset and results can be found in results section.

2 Related Work

Several approaches have been proposed for pairwise linkage classification using different machine learning algorithms. Han et al. [8] proposed two approaches using a Hybrid Naive Bayes and support vector machine(SVM) classifier. Huang et al. [11] used an online active SVM(LASVM) to boost the speed of SVM classifier. Song et al. [18] used probabilistic latent semantic analysis(pLSA) and Latent Dirichlet allocation(LDA) to disambiguate names based on publication

[1] List of common Chinese surnames, in Wikipedia. `https://en.wikipedia.org/wiki/List_of_common_Chinese_surname`

[2] `http://www.dev.patentsview.org/workshop`

content. Treeratpituk and Giles [21] first introduced the random forest (RF) for disambiguation and showed the random forest classifier at the time to have the best accuracy compared to other machine learning based classifiers. Godoi et al. [7] used an iterative approach to update ambiguous linkage with user feedback. Fan et al. [5] used graph based framework for name disambiguation, and Hermansson et al. [10] used graph kernels to calculate similarity based on local neighborhood structure. Instead of using machine learning algorithms, Santana et al. [17] used domain specific heuristics for classification. Recently, Ventura et al. [22] applied a random forest classifier with agglomerative clustering for inventor name disambiguation for USPTO database.

For clustering algorithms for disambiguation, Mann and Yarowsky [16] used a simple agglomerative clustering which still had a transitivity problem. The transitivity problem occurs when there are three records a, b, c and while a matches to b, b matches to c, a does not match with c. Han et al. [9] used K-spectral clustering which had scaling issues and the K (number of clusters) was heuristically determined. To overcome those problems, Huang et al. [11] proposed a density-based clustering(DBSCAN) algorithm. Another is a graphical approach using conditional random fields for clustering, using Markov Chain Monte Carlo(MCMC) methods [23]. Recently Khabsa et al. [13] proposed a constraint-based clustering algorithm based on DBSCAN and extended it to handle online clustering.

3 Disambiguation Process

Patent records have consistent metadata to that of scholarly publications. There exists title, personal information of inventors, such as name, affiliation, etc. Ventura et al. [22] applied author name disambiguation algorithms to patent records, showing very promising results. Our algorithm follows the same general steps of author name disambiguation.

First we train a pairwise classifier that determines whether each pair of inventor records is same person or not. Second, we apply blocking to the entire records for scaling. Finally, we cluster inventor records from each block separately using the classifier learned from the previous step.

3.1 Training Pairwise Classifier

Pairwise classifier is needed to distinguish whether each pair of inventor records is the same person or not. In this section we show what features are used and how we sample the training data. We compare several machine learning classifiers to find the best one for inventor name disambiguation.

3.1.1 Selecting Features

We start with the feature set used in Ventura et al. [22], and test additional features that are used in author disambiguation for scholarly databases. We

Category	Subcategory	Features
Inventor	First name	Exact, Jaro-Winkler, Soundex
	Middle name	Exact, Jaro-Winkler, Soundex
	Last name	Exact, Jaro-Winkler, Soundex, IDF
	Suffix	Exact
	Order	Order comparision
Affiliation	City	Exact, Jaro-Winkler, Soundex
	State	Exact
	Country	Exact
Co-author	Last name	# of name shared, IDF, Jaccard
Assignee	Last name	Exact, Jaro-Winkler, Soundex
Group	Group	Exact
	Subgroup	Exact
Title	Title	# of term shared

Table 1: Features used for the random forest

only kept features that had a meaningful decrease in Gini importance if they were removed. Table 1 shows all features used for the random forest classifier. A detailed explanation of each term is as follows:

- Exact: Exact string match, 3 if name matches and both full names, 2 if initial matches and not both full names, 1 if initial not matches and not both full names, 0 if name not matches and both full names.
- Jaro-Winkler: Jaro-Winkler distance [24] of two strings. Jaro-Winkler distance is a variant of Jaro distance. Jaro disatnace d_j of two string s_1 and s_2 is calculated as

$$d_j = \begin{cases} 0 & \text{if } n_m = 0 \\ \frac{1}{3}\left(\frac{n_m}{|s_1|} + \frac{n_m}{|s_2|} + \frac{n_m - \frac{1}{2}n_t}{n_m}\right) & \text{otherwise} \end{cases}$$

  where n_m is number of matching characters, and n_t is number of transpositions. Each character is considered as a match only if they are within distance of a half length of a longer string -1. Jaro-Winkler distance d_{jw} of two strings are calculated using this Jaro distance d_j,

$$d_{jw} = d_j + l_{\text{prefix}} p(1 - d_j)$$

  where l_{prefix} is length of common prefix between two strings (up to 4 characters). p is a scaling factor, we use 0.1.
- Soundex: Convert each string with Soundex algorithm [15] and then do an exact string match giving credit for phonetically similar strings. The basic idea of soundex algorithm is to cluster phonetically similar consonants and convert them with the group number. {b, f, p, v}, {c, g, j, k, q, s, x, z}, {d,t}, {l}, {m,n}, {r} are the 6 groups.
- IDF: Inverse document frequency(calculated by # of records total/# of records with name) of the name, to give more weight to a unique name.

Method	Precision	Recall	F1
Naive Bayes	0.9246	0.9527	0.9384
Logistic Regression	0.9481	0.9877	0.9470
SVM	0.9613	**0.9958**	0.9782
Decision Tree	0.9781	0.9798	0.9789
Conditional Inference Tree	0.9821	0.9879	0.9850
Random Forest	**0.9839**	0.9946	**0.9892**

Table 2: Comparison of different classification methods

- Order comparison: 2 if both records are first author, 1 if both records are last author, 0 otherwise.
- # of name shared: number of same name shared without considering order.
- # of term shared: number of shared terms appear in both titles, excluding common stop words.

3.1.2 Selecting Samples for Training Classifier

The existing labeled data from the USPTO database has two challenges in that we cannot directly use all possible pairs as a training set for a classifier. First, the majority of the labeled clusters have only a single record. In the Mixture dataset, these are 3,491 clusters out of 4,956 clusters(70.44%) and in the Common characteristics dataset 26,648 clusters out of 30,745 clusters(86.67%) that have only a single record. Those clusters are not useful as training data because we can only get negative pairs (two records not from the same person) from them. To train a good classifier, we need data that can give both positive and negative examples. As such we removed all those clusters, and used clusters that only have more than 1 record.

Second, there were insufficient informative negative samples from the labeled datasets. Since we need to use blocking for scaling, we want to use only pairs that consist of records from same block, since pairs from different blocks are not going to be examined in the clustering process. But there were few different clusters within each block in the labeled datasets. Since there were fewer negative samples than positive samples and to avoid overfitting, we take samples from a bigger block than the actual blocking size, using *first 3 characters of last name+first name initial* while actual blocking is done with *full last name+first name initial*.

3.1.3 Classifier Selection

We experimented with several supervised classifiers using the proposed feature set. We tested with the mixture of two training datasets - Mixture and Common characteristics. A detailed explanation of the datasets are in the result

Rank	Feature
1	Last name (Jaro-Winkler)
2	First name (Jaro-Winkler)
3	Last name (Exact)
4	Last name (Soundex)
5	First name (Soundex)
6	Affiliation (Jaro-Winkler)
7	First name (Exact)
8	State (Exact)
9	Middle name (Soundex)
10	Middle name (Exact)

Table 3: Top 10 important features of the random forest with respect to the Gini decrease

section. Table 2 shows the results with 4-fold cross validation. Tree-based classifiers have a higher accuracy compared to non-tree classifiers such as SVM and logistic regression. Among them, the Random Forest (RF) classifier gives the best accuracy in terms of F1 score. RF is an ensemble classifier that aggregates the votes from decision trees for classification [3]. Previous work [21] showed RF is effective for pairwise linkage classification in scholarly database. This experiments showed that for patent database RF also has the best accuracy. We trained our RF classifier with 100 trees with 5 features tried in each split. We estimated the out-of-bag (OOB) error of the RF to measure the classification quality. OOB error is known to be an unbiased estimation of test set classification error [3]. The error rates for Common characteristics and Mixture dataset were 0.05% and 0.07% respectively.

3.2 Blocking

The USPTO patent database consists of 12 millions of inventor mentions. Due to the limitation of physical memory, we cannot efficiently perform the clustering for the whole database. Blocking is done in preprocessing in order to solve this problem . Records are split into several blocks, based on the blocking function. The function should be carefully selected so that records from the same person are in the same block with high probability [2]. Then we perform clustering for each block.

Table 3 shows the top 10 important features of the RF according to the average Gini decrease. The table shows the most important features are from first name and last name. All features from them are in top 10 except for the last name IDF. Thus, for best performance we made a blocking function with a combination of the first name and last name. Figure 1 and Table 4 shows the accuracy and computation time with respect to different block sizes. While precision is steady with different block sizes, recall gets lower as the block size becomes smaller, as does F1. This is because this blocking function splits the potential matches into different blocks. While the accuracy is getting lower, the computation time is reduced due to smaller block sizes. We use *full last*

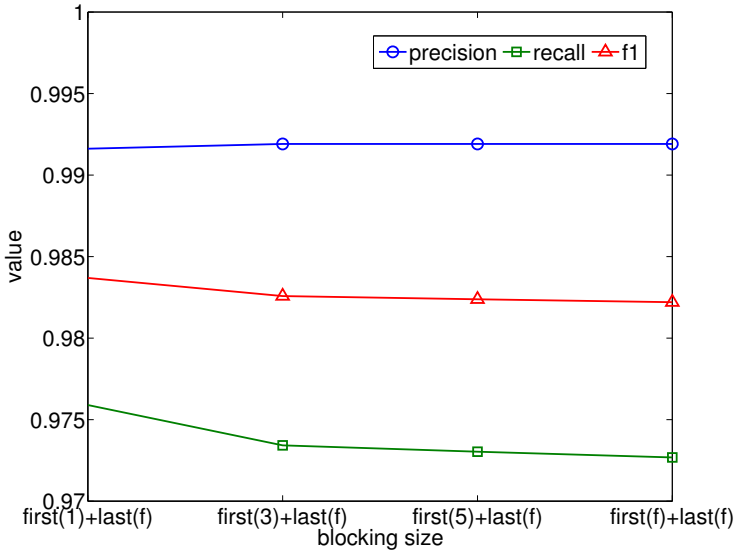

Figure 1: Evaluation of different blocking size. FN denotes the first name and LN denotes the last name. The number (n) denotes first n characters used for blocking. If n is f, full name was used.

Block	FN(1)+LN(f)	FN(3)+LN(f)	FN(5)+LN(f)	FN(f)+LN(f)
Time	6h 30min	5h 49min	5h 27min	5h 17min

Table 4: Computation time comparision for different block size

name+initial of first name, which was the blocking function that gives highest accuracy and that each block can be loaded fully into memory.

3.3 Clustering Using DBSCAN

We use a density-based clustering algorithm, DBSCAN [4] to cluster inventor records. DBSCAN is widely used for disambiguation, because it does not require a prior the number of clusters, and it resolves the transitivity problem [11].

Using DBSCAN to cluster inventor records, we need to define a distance function for each pair of inventor records. The RF classifier predicts whether each pair of records are from the same person or not with a binary value(0 or 1) output. From the RF, we can get the number of negative/positive votes in its trees. We use the fraction of negative(0) votes of the trees in random forest as the distance function [21]. The final resulting clusters from DBSCAN algorithm are the result of the disambiguation.

3.4 Parallelization

We use parallelization of GNU Parallel [19] to utilize all cores available for clustering [12]. Our work consumes memory proportional to the total number of records in the block. Due to the limitation of our physical memory, we cannot completely utilize all cores at a time if block size is too large. As such, we grouped all blocks with respect to total number of records.

The machine we use for the experiment has about 40GB memory available, and 12 cores(runs up to 24 threads simultaneously) at best. The first group consists of blocks that have less than 500 records, and we run 24 threads maximum simultaneously. The second group consists of blocks that have between 500 and 5,000 records and we run 12 threads maximum. The last group consists of blocks that have more than 5,000 records and we run 6 threads maximum.

4 Results

We tested our algorithm on the USPTO patent database. We used the same evaluation datasets of USPTO Patentsview inventor named disambiguation competition to compare the results. The test dataset includes ALS, ALS common, IS, E&S, and Phase2. The ALS and ALS common datasets are from Azoulay et al. [1], which consists of inventors from the Association of Medical Colleges(AAMC) Faculty Roster. ALS common is a subset of the ALS dataset with common popular names. The IS dataset is from Trajtenberg and Shiff [20], containing Israeli inventors in USPTO database. E&S dataset is from Ge et al. [6] and consists of patents from engineers and scientists. Phase2 is a random mixture of previous datasets. The training dataset includes the Mixture and Common characteristics datasets. Mixture dataset is random mixture of IS and E&S dataset, and Common characteristics dataset is a subsample of E&S dataset which was subsampled according to the match characteristics of the USPTO database, in terms of the mean number of inventors per patent and percentage of missing assignees.

The result of our disambiguated USPTO database shows a similar tendency to previous disambiguation studies of scholarly databases. Figure 2 shows the cluster frequency of each cluster size of CiteSeerX[3] and USPTO database after disambiguation. For both databases, small clusters have high frequency and big clusters are rare with a long tail. A total of 1.11 million clusters are produced by inventor name disambiguation and the average number of patent mentions per individual inventor is 4.93. For CiteSeerX database, the average number is 6.07.

For further evaluation, we measured pairwise precision, recall, and F1 score with definitions:

$$\text{Pairwise Precision} = \frac{\text{\# of correctly matched pairs}}{\text{\# of all matched pairs by algorithm}}$$

[3] `http://citeseerx.ist.psu.edu`

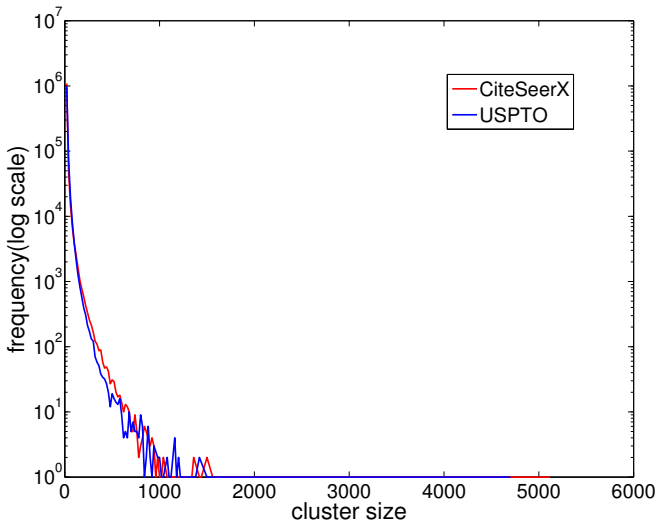

Figure 2: Frequency of each cluster size for the CiteSeerX and USPTO database

$$\text{Pairwise Recall} = \frac{\#\text{ of correctly matched pairs}}{\#\text{ of pairs in manually labeled dataset}}$$

$$\text{Pairwise F1 Score} = 2 \cdot \frac{\text{Pairwise Precision} \cdot \text{Pairwise Recall}}{\text{Pairwise Precision} + \text{Pairwise Recall}}$$

Table 5 shows the results for each training and test dataset. Results were slightly better with the Common characteristics dataset, as expected from OOB error of RF. This is because common characteristics dataset has more samples and is subsampled according to match the characteristics of the whole USPTO database. We can also see that the recall is relatively lower compare to the precision. Blocking affects the recall, as it can remove some potential matches. Since we have a trade-off between efficiency and recall in our algorithm, blocking needs to be further improved for higher recall. Table 6 shows F1 score comparison between our work and the best result from the competition for each test dataset. The winner of the competition used a pre-defined distance metric and Markov Chain Monte Carlo(MCMC) based clustering method inspired from [23]. Note that our algorithm has the best performance on all datasets. The P value with one-tailed Wilcoxon test is 0.03125, which indicates that the improvement of our algorithm is statistically significant at the 0.05 level. We can see from the results that the DBSCAN algorithm with the RF classifier used in scholarly disambiguation is also effective for inventor name disambiguation in a patent database.

Our disambiguation is much faster with parallelization. We used Intel Xeon

Test Set	Training Set	Precision	Recall	F1 Score
ALS	Mixture	0.9963	0.9790	0.9786
	Common	0.9960	0.9848	0.9904
ALS common	Mixture	0.9841	0.9796	0.9818
	Common	0.9820	0.9916	0.9868
IS	Mixture	0.9989	0.9813	0.9900
	Common	0.9989	0.9813	0.9900
E&S	Mixture	0.9992	0.9805	0.9898
	Common	0.9995	0.9810	0.9902
Phase2	Mixture	0.9912	0.9760	0.9836
	Common	0.9916	0.9759	0.9837

Table 5: Disambiguation evaluation

Test Set	F1(Ours)	F1(Winner)
ALS	0.9904	0.9879
ALS common	0.9868	0.9815
IS	0.9900	0.9783
E&S	0.9902	0.9835
Phase2	0.9837	0.9826
Average(±stddev.)	0.9882±0.0029	0.9827±0.0035

Table 6: Comparison with the competition winner

X5660@2.80GHz machine with 12 cores and 40GB memory available in an idle state, configured with RHEL 6. The disambiguation process takes about 6.5 hours to finish for both training sets. Currently we cannot fully utilize all the CPUs for certain blocks that contain large number of records, because of memory limitations. Better way of blocking such as [2] is needed for efficient memory usage, for fast performance and scalability. This remains as a future work.

5 Conclusions

We present a machine learning based algorithm for inventor name disambiguation for patent database. Motivated by the feature set of author name disambiguation for scholarly databases, we devised a proposed feature set that showed a significant low OOB error rate, 0.05% at minimum. Based on experiments with several machine learning classifiers, we use random forest classifier to determine whether each pair of inventor records are from a the same inventor or not. Disambiguation is done by using DBSCAN clustering algorithm. We define distance function of each pair of inventor records as the ratio of votes in random forest classifier. In addition to make the algorithm scalable, we use blocking and parallelization, scheduling threads based on the size of blocks. Evaluation results with the dataset from USPTO PatentsView inventor name disambiguation competition shows our algorithm outperforms all algorithms submitted to the

competition in comparable running time.

Currently our algorithm is memory bounded since a great deal of memory is used to store inventor information needed for calculating features. This becomes a bottleneck when we parallelize the algorithm, since some of the block is huge due to the popularity of certain names. In future work, one could explore a better method for blocking for efficient memory usage. It would be interesting to see if other methods using graph or link data could be incorporated for better performance.

6 Acknowledgments

We gratefully acknowledge Evgeny Klochikhin and Ahmad Emad for assistance in the evaluation of the dataset used in USPTO PatentsView inventor name disambiguation competition and partial support from the National Science Foundation.

References

[1] P. Azoulay, R. Michigan, and B. N. Sampat. The anatomy of medical school patenting. *New England Journal of Medicine*, 357(20):2049–2056, 2007.

[2] M. Bilenko, B. Kamath, and R. J. Mooney. Adaptive blocking: Learning to scale up record linkage. In *Proceedings of the 6th IEEE International Conference on Data Mining(ICDM'06)*, pages 87–96, 2006.

[3] L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

[4] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD'96)*, volume 96, pages 226–231, 1996.

[5] X. Fan, J. Wang, X. Pu, L. Zhou, and B. Lv. On graph-based name disambiguation. *Journal of Data and Information Quality (JDIQ)*, 2(2):10, 2011.

[6] C. Ge, K.-W. Huang, and I. Png. Engineer/scientist careers: Patents, online profiles, and misclassification bias. *Strategic Management Journal*, 37(1):232–253, 2016.

[7] T. A. Godoi, R. d. S. Torres, A. M. Carvalho, A. A. Gonçalves, Marcos A .and Ferreira, W. Fan, and E. A. Fox. A relevance feedback approach for the author name disambiguation problem. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries(JCDL'13)*, pages 209–218, 2013.

[8] H. Han, C. L. Giles, H. Zha, C. Li, and K. Tsioutsiouliklis. Two supervised learning approaches for name disambiguation in author citations. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries(JCDL'04)*, pages 296–305, 2004.

[9] H. Han, H. Zha, and C. L. Giles. Name disambiguation in author citations using a k-way spectral clustering method. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries(JCDL'05)*, pages 334–343, 2005.

[10] L. Hermansson, T. Kerola, F. Johansson, V. Jethava, and D. Dubhashi. Entity disambiguation in anonymized graphs using graph kernels. In *Proceedings of the 22nd ACM International Conference on information & knowledge management(CIKM'13)*, pages 1037–1046, 2013.

[11] J. Huang, S. Ertekin, and C. L. Giles. Efficient name disambiguation for large-scale databases. In *Proceedings of the 10th European Conference on Principle and Practice of Knowledge Discovery in Databases(PKDD'06)*, pages 536–544, 2006.

[12] M. Khabsa, P. Treeratpituk, and C. L. Giles. Large scale author name disambiguation in digital libraries. In *IEEE International Conference on Big Data*, pages 41–42, 2014.

[13] M. Khabsa, P. Treeratpituk, and C. L. Giles. Online person name disambiguation with constraints. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries(JCDL'15)*, pages 37–46, 2015.

[14] K. Kim, M. Khabsa, and C. L. Giles. Inventor name disambiguation for a patent database using a random forest and dbscan. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries(JCDL'16)*, pages 269–270, 2016.

[15] D. E. Knuth. *The Art of Computer Programming: Sorting and Searching*, volume 3. Addison-Wesley, 1973.

[16] G. S. Mann and D. Yarowsky. Unsupervised personal name disambiguation. In *Proceedings of the seventh conference on Natural language learning at HLT-NAACL 2003*, volume 4, pages 33–40. Association for Computational Linguistics, 2003.

[17] A. F. Santana, M. A. Gonçalves, A. H. Laender, and A. Ferreira. Combining domain-specific heuristics for author name disambiguation. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries(JCDL'14)*, pages 173–182, 2014.

[18] Y. Song, J. Huang, I. G. Councill, J. Li, and C. L. Giles. Efficient topic-based unsupervised name disambiguation. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries(JCDL'07)*, pages 342–351, 2007.

[19] O. Tange et al. Gnu parallel-the command-line power tool. *The USENIX Magazine*, 36(1):42–47, 2011.

[20] M. Trajtenberg and G. Shiff. *Identification and mobility of Israeli patenting inventors*. Pinhas Sapir Center for Development, 2008.

[21] P. Treeratpituk and C. L. Giles. Disambiguating authors in academic publications using random forests. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries(JCDL'09)*, pages 39–48, 2009.

[22] S. L. Ventura, R. Nugent, and E. R. Fuchs. Seeing the non-stars:(some) sources of bias in past disambiguation approaches and a new public tool leveraging labeled records. *Research Policy*, 2015.

[23] M. Wick, S. Singh, and A. McCallum. A discriminative hierarchical model for fast coreference at large scale. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics(ACL'12)*, pages 379–388. Association for Computational Linguistics, 2012.

[24] W. E. Winkler. String comparator metrics and enhanced decision rules in the fellegi-sunter model of record linkage. In *Proceedings of the Section on Survey Research Methods*, pages 354–359. American Statistical Association, 1990.

[25] W. E. Winkler. Matching and record linkage. *Wiley Interdisciplinary Reviews: Computational Statistics*, 6(5):313–325, 2014.